Recognizing Union-Find trees is **NP**-complete[☆]

Kitti Gelle[1], Szabolcs Iván[1]

[a]*University of Szeged, Hungary*

Abstract

Disjoint-Set forests, consisting of Union-Find trees are data structures having a widespread practical application due to their efficiency. Despite them being well-known, no exact structural characterization of these trees is known (such a characterization exists for Union trees which are constructed without using path compression). In this paper we provide such a characterization and show that the decision problem whether a given tree is a Union-Find tree is **NP**-complete.

arXiv:1510.07462v1 [cs.CC] 26 Oct 2015

1. Introduction

Disjoint-Set forests, introduced in [4], are fundamental data structures in many practical algorithms where one has to maintain a partition of some set, which support three operations: *creating* a partition consisting of singletons, *querying* whether two given elements are in the same class of the partition (or equivalently: *finding* a representative of a class, given an element of it) and *merging* two classes. Practical examples include e.g. building a minimum-cost spanning tree of a weighted graph [2], unification algorithms [6] etc.

To support these operations, even a linked list representation suffices but to achieve an almost-constant amortized time cost per operation, Disjoint-Set forests are used in practice. In this data structure, sets are represented as directed trees with the edges directed towards the root; the CREATE operation creates n trees having one node each (here n stands for the number of the elements in the universe), the FIND operation takes a node and returns the root of the tree in which the node is present (thus the SAME-CLASS(x,y) operation is implemented as FIND(x) == FIND(y)), and the MERGE(x,y) operation is implemented by merging the trees containing x and y, i.e. making one of the root nodes to be a child of the other root node (if the two nodes are in different classes).

In order to achieve near-constant efficiency, one has to keep the (average) height of the trees small. There are two "orthogonal" methods to do that: first, during the merge operation it is advisable to attach the "smaller" tree below the "larger" one. If the "size" of a tree is the number of its nodes, we say the trees are built up according to the *union-by-size* strategy, if it's the depth of a tree, then we talk about the *union-by-rank* strategy. Second, during a FIND operation invoked on some node x of a tree, one can apply the *path compression* method, which reattaches each ancestor of x directly to the root of the tree in which they are present. If one applies both the path compression method and either one of the union-by-size or union-by-rank strategies, then any sequence of m operations on a universe of n elements has worst-case time cost $O(m\alpha(n))$ where α is the inverse of the extremely fast growing (not primitive recursive) Ackermann function for which $\alpha(n) \leq 5$ for each practical value of n (say, below 2^{65535}), hence it has an amortized almost-constant time cost [7]. Since it's proven [3] that *any* data structure has worst-case time cost $\Omega(m\alpha(n))$, the Disjoint-Set forests equipped with a strategy and path compression offer a theoretically optimal data structure which performs exceptionally well also in practice, not only in theory. For more details see standard textbooks on data structures, e.g. [2].

Due to these facts, it is certainly interesting both from the theoretical as well as the practical point of view to characterize those trees that can arise from a forest of singletons after a number of merging and finding operations, which we call Union-Find trees in this paper. One could e.g. test Disjoint-Set implementations since if at any given point of execution a tree of a Disjoint-Set forest is not a valid Union-Find tree, then it is certain that there is a bug in the implementation of the data structure. Nevertheless, only the characterization of Union trees is known up till now [1], i.e. which correspond to the case when one uses one of the union-by-x strategies but *not* path compression. Since in that case the data structure offers only a theoretic bound of $\Theta(\log n)$ on the amortized time cost, in practice all implementations imbue path compression as well, so for a characterization to be really useful, it has to cover this case as well.

In this paper we show that the recognition problem of Union-Find trees is **NP**-complete when the union-by-size strategy is used (and leave open the case of the union-by-rank strategy), coupled with path compression, confirming the statement from [1] that this problem "seems to be much harder" than recognizing Union trees (which in turn can be done in low-degree polynomial time).

2. Notation

A *tree* is a tuple $t = (V_t, \text{ROOT}_t, \text{PARENT}_t)$ with V_t being the finite set of its *nodes*, $\text{ROOT}_t \in V_t$ its *root* and $\text{PARENT}_t : (V_t - \{\text{ROOT}_t\}) \to V_t$ mapping each nonroot node to its parent (so that the graph of

[☆]This research was supported by NKFI grant number K108448.

PARENT_t is a directed acyclic graph, or a DAG, with edges being directed towards the root).

For a tree t and a node $x \in V_t$, let $\text{CHILDREN}(t, x)$ stand for the set $\{y \in V_t : \text{PARENT}_t(y) = x\}$ of its children and $\text{CHILDREN}(t)$ stand as a shorthand for $\text{CHILDREN}(t, \text{ROOT}_t)$, the set of depth-one nodes of t. Also, let $x \preceq_t y$ denote that x is an *ancestor* of y in t, i.e. $x = \text{PARENT}_t^k(y)$ for some $k \geq 0$ and let $\text{SIZE}(t, x) = |\{y \in V_t : x \preceq_t y\}|$ stand for the number of *descendants* of x (including x itself). Let $\text{SIZE}(t)$ stand for $\text{SIZE}(t, \text{ROOT}_t)$, the number of nodes in the tree t. For $x \in V_t$, let $t|_x$ stand for the *subtree* $(V_x = \{y \in V_t : x \preceq_t y\}, x, \text{PARENT}_t|_{V_x})$ of t rooted at x.

Two operations on trees are that of *merging* and *collapsing*. Given two trees $t = (V_t, \text{ROOT}_t, \text{PARENT}_t)$ and $s = (V_s, \text{ROOT}_s, \text{PARENT}_s)$ with V_t and V_s being disjoint, their merge $\text{MERGE}(t, s)$ (in this order) is the tree $(V_t \cup V_s, \text{ROOT}_t, \text{PARENT})$ with $\text{PARENT}(x) = \text{PARENT}_t(x)$ for $x \in V_t$, $\text{PARENT}(\text{ROOT}_s) = \text{ROOT}_t$ and $\text{PARENT}(y) = \text{PARENT}_s(y)$ for each nonroot node $y \in V_s$ of s. Given a tree $t = (V, \text{ROOT}, \text{PARENT})$ and a node $x \in V$, then $\text{COLLAPSE}(t, x)$ is the tree $(V, \text{ROOT}, \text{PARENT}')$ with $\text{PARENT}'(y) = \text{ROOT}$ if y is a nonroot ancestor of x in t and $\text{PARENT}'(y) = \text{PARENT}(y)$ otherwise. For examples, see Figure 1.

The class of Union trees is the least class of trees satisfying the following two conditions: every singleton tree (having exactly one node) is a Union tree, and if t and s are Union trees with $\text{SIZE}(t) \geq \text{SIZE}(s)$, then $\text{MERGE}(t, s)$ is a Union tree as well.

Analogously, the class of Union-Find trees is the least class of trees satisfying the following three conditions: every singleton tree is a Union-Find tree, if t and s are Union-Find trees with $\text{SIZE}(t) \geq \text{SIZE}(s)$, then $\text{MERGE}(t, s)$ is a Union-Find tree as well, and if t is a Union-Find tree and $x \in V_t$ is a node of t, then $\text{COLLAPSE}(t, x)$ is also a Union-Find tree.

We'll frequently sum the size of "small enough" children of nodes, so we introduce a shorthand also for that: for a tree t, node x of t and threshold $W \geq 0$, let $\text{SUMSIZE}(t, x, W)$ stand for $\sum_{y \in \text{CHILDREN}(t,x), \text{SIZE}(t,y) \leq W} \text{SIZE}(t, y)$. We say that a node x of a tree t *satisfies the Union condition* if for each child y of x we have $\text{SUMSIZE}(t, x, W) \geq W$ where $W = \text{SIZE}(t, y) - 1$. Otherwise, we say that x *violates the Union condition* (at child y). Then, the characterization of Union trees from [1] can be formulated in our terms as follows:

Theorem 1. *A tree t is a Union tree if and only if for each node x of t satisfies the Union condition.*

Equivalently, x satisisfies the Union condition if $x_1, \ldots, x_k$ is an enumeration of $\text{CHILDREN}(t, x)$ such that $\text{SIZE}(t, x_i) \leq \text{SIZE}(t, x_{i+1})$ for each i, then $\sum_{j<i} \text{SIZE}(t, x_j) \geq \text{SIZE}(t, x_i) - 1$ for each $i = 1, \ldots, k$. (In particular, each non-leaf node has to have a leaf child.)

3. Structural characterization of Union-Find trees

Suppose s and t are trees on the same set V of nodes and with the same root r. We write $s \preceq t$ if $x \preceq_s y$ implies $x \preceq_t y$ for each $x, y \in V$. Clearly, $\preceq$ is a partial order on any set of trees (i.e. is a reflexive, transitive and antisymmetric relation). It is also clear that $s \preceq t$ if and only if $\text{PARENT}_s(x) \preceq_t x$ holds for each $x \in V - \{r\}$ which is further equivalent to requiring $\text{PARENT}_s(x) \preceq_t \text{PARENT}_t(x)$ since $\text{PARENT}_s(x)$ cannot be x.

Another notion we define is the (partial) operation PUSH on trees as follows: when t is a tree and $x \neq y \in V_t$ are *siblings* in t, i.e. have the same parent, then $\text{PUSH}(t, x, y)$ is defined as the tree $(V_t, \text{ROOT}_t, \text{PARENT}')$ with $\text{PARENT}'(z) = \begin{cases} y & \text{if } z = x, \\ \text{PARENT}_t(z) & \text{otherwise} \end{cases}$, that is, we "push" the node x one level deeper in the tree just below its former sibling y. (See Figure 1.)

We write $t \vdash t'$ when $t' = \text{PUSH}(t, x, y)$ for some x and y, and as usual, $\vdash^*$ denotes the reflexive-transitive closure of $\vdash$. Observe that when $t' = \text{PUSH}(t, x, y)$, then $\text{SIZE}(t', y) = \text{SIZE}(t, y) + \text{SIZE}(t, x) > \text{SIZE}(t, y)$ and $\text{SIZE}(t', z) = \text{SIZE}(t, z)$ for each $z \neq y$, hence $\vdash^*$ is also a partial ordering on trees. This is not a mere coincidence:

Proposition 1. *For any pair s and t of trees, $s \preceq t$ if and only if $s \vdash^* t$.*

PROOF. For $\vdash^*$ implying $\preceq$ it suffices to show that $\vdash$ implies $\preceq$ since the latter is a preorder. So let $s = (V, r, \text{PARENT})$ and $x \neq y \in V$ be siblings in s with the common parent z and let $t = \text{PUSH}(s, x, y)$. Then, since $\text{PARENT}_s(x) = z = \text{PARENT}_t(y) = \text{PARENT}_t(\text{PARENT}_t(x))$, we get $\text{PARENT}_s(x) \preceq_t x$, and by $\text{PARENT}_s(w) = \text{PARENT}_t(w)$ for each node $w \neq x$, we have $s \preceq t$. It is clear that $\preceq$ is equality on singleton trees, thus $\preceq$ implies $\vdash^*$ for trees of size 1. We apply induction on the size of $s = (V, r, \text{PARENT})$ to show that whenever $s \preceq t = (V, r, \text{PARENT}')$, then $s \vdash^* t$ as well. Let X stand for the set $\text{CHILDREN}(s)$ of depth-one nodes of s and Y stand for $\text{CHILDREN}(t)$. Clearly, $Y \subseteq X$ since by $s \preceq t$, any node x of s having depth at least two has to satisfy $\text{PARENT}(x) \preceq_t \text{PARENT}'(x)$ and since $\text{PARENT}(x) \neq r$ for such nodes, x has to have depth at least two in t as well. Now for any node $x \in X - Y$ there exists a unique node $y(x) \in Y$ such that $y(x) \preceq_t x$. Let $\{x_1, \ldots, x_k\}$ stand for $X - Y$ and let us define the trees $s_0 = s$, $s_i = \text{PUSH}(s_{i-1}, x_i, y(x_i))$. Then $s \vdash^* s_k$, $\text{CHILDREN}(s_k) = Y = \text{CHILDREN}(t)$ and for each $y \in Y$, $s_k|_y \preceq t|_y$. Applying the induction hypothesis we get that $s_k|_y \vdash^* t|_y$ for each $y \in Y$, hence the immediate subtrees of s_k can be transformed into the immediate subtrees of t by repeatedly applying PUSH operations, hence $s \vdash^* t$ as well.

The relations $\preceq$ and $\vdash^*$ are introduced due to their intimate relation to Union-Find and Union trees:

Theorem 2. *A tree t is a Union-Find tree if and only if $t \vdash^* s$ for some Union tree s.*

PROOF. Let t be a Union-Find tree. We show the claim by structural induction. For singleton trees the claim holds since any singleton tree is a Union tree as well. Suppose $t = \text{MERGE}(t_1, t_2)$. Then by the induction hypothesis, $t_1 \vdash^* s_1$ and $t_2 \vdash^* s_2$ for the Union trees s_1 and s_2. Then, for the tree $s = \text{MERGE}(s_1, s_2)$ we get that $t \vdash^* s$. Finally, assume $t = \text{COLLAPSE}(t', x)$ for some node x. Let $x = x_1 \succ x_2 \succ \ldots \succ x_k = \text{ROOT}_{t'}$ be the ancestral sequence of x in t'. Then, defining $t_0 = t$, $t_i = \text{PUSH}(t_{i-1}, x_i, x_{i+1})$ we get that $t \vdash^* t_{k-1} = t'$ and $t' \vdash^* s$ for some Union tree applying the induction hypothesis, thus $t \vdash^* s$ also holds.

Now assume $t \vdash^* s$ (equivalently, $t \preceq s$) for some Union tree s. We show the claim by induction on $\text{SIZE}(t)$. For singleton trees the claim holds since any singleton tree is a Union-Find tree. Let $t = (V, r, \text{PARENT})$ be a tree with $t \preceq s = (V, r, \text{PARENT}')$ for the Union tree s and assume the claim holds for any tree having less than $|V|$ nodes. Let X stand for the set $\text{CHILDREN}(t)$ of depth-one nodes of t and Y stand for $\text{CHILDREN}(s)$. By $t \preceq s$ we get that $Y \subseteq X$. Let us define for each $y \in Y$ the tree $t_y = (V_y, y, \text{PARENT}_y)$ with $V_y = \{x \in V : y \preceq_s x\}$ and $\text{PARENT}_y(x) = \begin{cases} y & \text{if } y \in V_y \cap X \\ \text{PARENT}_t(x) & \text{otherwise} \end{cases}$ for each $x \in V_y - \{y\}$. Then, $t_y \preceq s|_y$ for each $y \in Y$, thus by the induction hypothesis, each tree t_y is a Union-Find tree. Since s is a Union tree, its root satisfy the Union condition, hence the members of $Y = \text{CHILDREN}(s)$ can be ordered as $Y = \{y_1, \ldots, y_k\}$ such that $\sum_{i=1}^{j-1} \text{SIZE}(s, y_i) \geq \text{SIZE}(s, y_j) - 1$ for each $j = 1, \ldots, k$. Thus, the tree t_0 being a singleton tree on r, and $t_i = \text{MERGE}(t_{i-1}, t_{y_i})$ is a valid Union-Find tree as well (by each t_y, $y \in Y$ being a Union-Find tree and by the Union condition on their respective sizes, since $\text{SIZE}(s, y) = \text{SIZE}(t_y)$ for each $y \in Y$). Finally, the tree t results from the tree t_k constructed above by applying successively one COLLAPSE operation on each node in $X - Y$, hence t is a Union-Find tree as well.

4. Complexity

In this section we show that recognition of Union-Find trees is **NP**-complete. (Note that membership in **NP** is clear by Theorem 2 and that the possible number of pushes is bounded above by n^2.)

Let $H > 0$ be a parameter (which we call the "heaviness threshold parameter" later on) to be specified later and $t = (V, \text{ROOT}, \text{PARENT})$ a tree. We call a node x *light* if $\text{SIZE}(t, x) \leq H$, *heavy* if $\text{SIZE}(t, x) > H$ and a *basket* if it has a heavy child in t. Note that every basket is heavy. A particular type of baskets is an *empty basket* which is a node having a single child which is a node of size $H + 1$.

Let us call a tree t_0 *flat* if it satisfies all the following conditions:

1. There are $K > 0$ depth-one nodes of t_0 which are empty baskets.
2. There is exactly one non-basket heavy depth-one node of t_0, with size $H + 1$.
3. The total size of light depth-one nodes is $(K + 1) \cdot H$.
4. The light nodes and non-basket heavy nodes are Union trees, e.g. have only direct children as descendants.

See Figure 2.

The fact that the PUSH operation cannot decrease the size of a node has a useful corollary:

Proposition 2. *If $t \vdash^* t'$, and x is a heavy (basket, resp.) node of t, then x is a heavy (basket, resp.) node of t' as well. Consequently, if x is light in t', then x is light in t as well.*

PROOF. Retaining heaviness simply comes from $\text{SIZE}(t, x) \leq \text{SIZE}(t', x)$. When x is a basket node of t, say having a heavy child y, then by $x \preceq_t y$ and $t \vdash^* t'$ we get $x \preceq_{t'} y$ as well, hence x has a (unique) child z which is an ancestor of y in t', thus by $\text{SIZE}(t', z) \geq \text{SIZE}(t', y) \geq \text{SIZE}(t, y)$ we get that z is heavy in t' hence x is a basket of t' as well.

We will frequently sum the sizes of the light children of nodes, so for a tree t and node x of t, let $\text{SUMLIGHT}(t, x)$ stand for $\text{SUMSIZE}(t, x, H)$, the total size of the light children of x in t.

We introduce a charge function c as follows: for a tree t and node $x \in V_t$, let the charge of x in t, denoted $c(t, x)$ be

$$c(t, x) = \begin{cases} 0 & \text{if } x \text{ is a light node of } t, \\ \text{SUMLIGHT}(t, x) - H & \text{otherwise.} \end{cases}$$

It's worth observing that when x is a non-basket heavy node of t, then $\text{SUMLIGHT}(t, x) = \text{SIZE}(t, x) - 1$ (since by being non-basket, each child of x is light).

Note that in particular the charge of a node is computable from the multiset of the sizes of its children (or, from the multiset of the sizes of its *light* children along with its own size). Let $t' = \text{PUSH}(t, x, y)$ and let $z = \text{PARENT}_t(y)$ be the common parent of x and y in t. Then, $c(t, w) = c(t', w)$ for each node $w \notin \{y, z\}$ since for those nodes this multiset does not change. We now list how $c(t', y)$ and $c(t', z)$ vary depending on whether y and x are light or heavy nodes. Note that $\text{SIZE}(t', z) = \text{SIZE}(t, z)$ (thus the charge of z can differ in t and t' only if z is heavy and $\text{SUMLIGHT}(t', z) \neq \text{SUMLIGHT}(t, z)$) and $\text{SIZE}(t', y) = \text{SIZE}(t, x) + \text{SIZE}(t, y)$.

i) If x and y are both light nodes with $\text{SIZE}(t, x) + \text{SIZE}(t, y) \leq H$ (i.e. y remains light in t' after x is pushed below y) then $\text{SUMLIGHT}(t', z) = \text{SUMLIGHT}(t, z)$ hence $c(t, z) = c(t', z)$ and $c(t', y) = c(t, y) = 0$ since y is still light in t'.

ii) If x and y are both light nodes with $\text{SIZE}(t, x) + \text{SIZE}(t, y) > H$ (i.e. y becomes heavy due to the pushing), then z is heavy as well (both in t and in t' since $\text{SIZE}(t, z) = \text{SIZE}(t', z)$). Then since $\text{SUMLIGHT}(t', z) = \text{SUMLIGHT}(t, z) - (\text{SIZE}(t, x) + \text{SIZE}(t, y))$, we get $c(t', z) = c(t, z) - (\text{SIZE}(t, x) + \text{SIZE}(t, y))$ and $c(t', y) = \text{SIZE}(t, y) + \text{SIZE}(t, x) - (H + 1)$ since y is heavy in t' having only light children (by the assumption, x is still light and each child of y already present in t have to be also light since y itself is light in t).

iii) If x is heavy and y is in light in t, then after pushing, y becomes heavy as well (and z is also heavy). Then, by $\text{SUMLIGHT}(t', z) = \text{SUMLIGHT}(t, z) - \text{SIZE}(t, y)$ (since the child y of z loses its status being light) we have $c(t', z) = c(t, z) - \text{SIZE}(t, y)$ and in t', y is a heavy node with light children of total size $\text{SIZE}(t, y) - 1$, hence $c(t', y) = \text{SIZE}(t, y) - (H + 1)$ (while $c(t, y) = 0$ since y is light in t).

iv) If x is light and y is heavy in t, then z is heavy as well, z loses its light child x while y gains the very same light child, hence $c(t', z) = c(t, z) - \text{SIZE}(t, x)$ and $c(t', y) = c(t, y) + \text{SIZE}(t, x)$.

v) Finally, if both x and y are heavy, then z is heavy as well, and since neither z nor y loses or gains a light child, $c(t', z) = c(t, z)$ and $c(t', y) = c(t, y)$.

Observe that in Cases i), iv) and v) we have $\sum_{x \in V_t} c(t, x) = \sum_{x \in V_t} c(t', x)$ while in Cases ii) and iii) it holds that $\sum_{x \in V_t} c(t, x) >$

$\sum_{x \in V_t} c(t', x)$ (namely, the total charge decreases by $H+1$ in these two cases).

Now for a flat tree t_0 it is easy to check that $\sum_{x \in V_t} c(t, x)$ is zero: we have $\textsc{sumlight}(t_0, \textsc{root}_{t_0}) = (K+1) \cdot H$ by assumption, hence the root node is in particular heavy and $c(t_0, \textsc{root}_{t_0}) = K \cdot H$, the empty baskets (having no light node at all) have charge $-H$ each and there are K of them, finally, all the heavy nodes of size $H+1$ have charge zero as well as each light node, making up the total sum of charges to be zero.

The charge function we introduced turns out to be particularly useful due to the following fact:

Proposition 3. *If t is a Union tree, then $c(t, x) \geq 0$ for each node of t.*

Proof. If x is a light node of t, then $c(t, x) = 0$ and we are done. If x is a heavy node which is *not* a basket, then all of its children are light nodes, thus $\textsc{sumlight}(t, x) = \textsc{size}(t, x) - 1$, making $c(t, x) = \textsc{size}(t, x) - (H+1)$ which is nonnegative since $\textsc{size}(t, x) > H$ by x being heavy.

Finally, if x is a basket node of t, then it has at least one heavy child. Let y be the smallest heavy child of x. Since t is a Union tree, we have $\textsc{sumsize}(t, x, \textsc{size}(t, y) - 1) \geq \textsc{size}(t, y) - 1$. By the choice of y, every child of x which is lighter than y is light itself, thus $\textsc{sumsize}(t, x, \textsc{size}(t, y) - 1) = \textsc{sumsize}(t, x, H) = \textsc{sumlight}(t, x)$, moreover, $\textsc{size}(t, y) - 1 \geq H$ since y is heavy, thus $\textsc{sumlight}(t, x) \geq H$ yielding $c(t, x) = \textsc{sumlight}(t, x) - H \geq 0$ and the statement is proved.

Thus, since a flat tree has total charge 0 while in a Union tree each node has nonnegative charge, and the push operation either decreases the total charge (in Cases ii) and iii) above) or leaves the total charge unchanged (in the other cases), we get the following:

Lemma. Suppose $t_0 \vdash^* t$ for a flat tree t_0 and a Union tree t. Then for any sequence of push operations transforming t_0 into t, each push operation is of type i), iv) or v) above, i.e.

- either a light node is pushed into a light node, yielding still a light node;
- or some node (either heavy or light) is pushed into a heavy node.

Moreover, the charge of each node of t has to be zero.

In particular, a heavy node x has charge 0 in a tree t iff $\textsc{sumlight}(t, x) = H$ yielding also that each non-basket heavy node of t has to be of size exactly $H+1$. It's also worth observing that by applying the above three operations one does not *create* a new heavy node: light nodes remain light all the time. Moreover, if for a basket node $\textsc{sumlight}(t, x) = H$ in a Union tree t, then there has to exist a heavy child of x in t of size exactly $H+1$ (by x being a basket node, there exists a heavy child and if the size of the lightest heavy child is more than $H+1$, then the Union condition gets violated).

Recall that in a flat tree t_0 with K empty baskets, there are $K+1$ baskets in total (the depth-one empty baskets and the root node), and there are $K+1$ non-basket heavy nodes (one in each basket, initially), each having size $H+1$, and all the other nodes are light.

Thus if $t_0 \vdash^* t$ for some Union tree t, then the set of non-basket heavy nodes coincide in t_0 and in t, and also in t, the size of each such node is still $H+1$. In particular, one cannot increase the size of such a node by pushing anything into such a node.

Summing up the result of the above reasoning we have:

Lemma. Suppose $t_0 \vdash^* t$ for some flat tree t_0 and Union tree t. Then for any pushing sequence transforming t_0 into t, each push has to have one of the following two forms:

- pushing a light node into a light node, the result remaining still light;
- pushing a node into a basket.

We can even restrict the order in which the above operations are applied. Indeed, assume $t' = \textsc{push}(t, x, y)$ for the light nodes x and y of t with $\textsc{size}(t, x) + \textsc{size}(t, y) \leq H$ and $t'' = \textsc{push}(t', z, w)$ for the basket w of t'. Then,

- if $z = y$, then we can get t'' from t by pushing x and y into the basket w first, then pushing x into y;
- otherwise we can simply swap the two push operations since w cannot be either x or y (since light nodes are not baskets), nor descendants of x or y, thus z and w are already siblings in t as well, hence z can be pushed into w also in t, and afterwards since x and y are siblings in the resulting tree, x can be pushed into y.

Thus if $t_0 \vdash^* t$ for some flat tree t_0 and Union tree t, we can assume that first we push nodes into baskets, then light nodes into light nodes, yielding light nodes. However, it turns out the latter type of pushing cannot fix the Union status of the trees we consider.

Proposition 4. *Suppose t is a tree with a basket node x violating the Union condition, i.e. for some child y of x it holds that $\textsc{sumsize}(t, x, \textsc{size}(t, y) - 1) < \textsc{size}(t, y) - 1$. Then for any tree $t' = \textsc{push}(t, z, w)$ with z and w being light nodes with total size at most H we have that x still violates the Union condition in t'.*

Proof. If z (and w) are not children of x, then $\textsc{children}(t, x) = \textsc{children}(t', x)$ (since in particular $w \neq x$ by w is light and x is a basket) and each child of x also has the same size in t' as in t, hence the Union condition is still violated. Now assume z and w are children of x. Upon pushing, x loses a child of size $\textsc{size}(t, z)$ and a child of size $\textsc{size}(t, w)$ and gains a child of size $\textsc{size}(t, z) + \textsc{size}(t, w)$. It is clear that $\textsc{sumsize}(t, x, W) \geq \textsc{sumsize}(t', x, W)$ for each possible W then (equality holds when $W \geq \textsc{size}(t, z) + \textsc{size}(t, w)$ or $W < \min\{\textsc{size}(t, w), \textsc{size}(t, z)\}$ and strict inequality holds in all the other cases). It is also clear that there is at least one child y' of x in t' such that $\textsc{size}(t, y) \leq \textsc{size}(t', y')$: if $y \neq z$ then $y' = y$ suffices, otherwise $y' = w$ is fine. Now let y_0 be such a child with $W = \textsc{size}(t', y_0)$ being the minimum possible. Then $\textsc{sumsize}(t', x, W - 1) = \textsc{sumsize}(t', x, \textsc{size}(t, y) - 1) \leq \textsc{sumsize}(t, x, \textsc{size}(t, y) - 1) < \textsc{size}(t, y) - 1 \leq \textsc{size}(t', y_0) - 1$ and thus x still violates the Union condition in t' as well.

Now observe that initially in any flat tree t_0, only basket nodes violate the Union condition. Moreover, by pushing arbitrary nodes into baskets only the baskets' Union status can change (namely, upon pushing into some basket x, the status of x and its parent can change, which is also a basket). Thus, after pushing nodes into baskets we either already have a Union tree, or not, but in the latter case we cannot transform the tree into a Union tree by pushing light nodes into light nodes. Hence we arrive to the following characterization:

Proposition 5. *Assume t_0 is a flat tree having K empty baskets. Then, t_0 is a Union-Find tree if and only if the set $\mathcal{L}$ of its depth-one light nodes can be partitioned into sets $\mathcal{L}_1, \ldots, \mathcal{L}_{K+1}$ such that for each $1 \le i \le K+1$, $\sum_{x \in \mathcal{L}_i} \textsc{size}(t_0, x) = H$ and for each $y \in \mathcal{L}_i$, $\sum_{z \in \mathcal{L}_i, \textsc{size}(t_0,z) < \textsc{size}(t_0,y)} \textsc{size}(t_0, z) \ge \textsc{size}(t_0, y) - 1$.*

Proof. Assume the set $\mathcal{L}$ of the depth-one light nodes of t_0 can be partitioned into sets $\mathcal{L}_i$, $i = 1, \ldots, K+1$ as above. Let $y_1, \ldots, y_K$ be the empty basket nodes of t_0. Then by pushing every member of $\mathcal{L}_i$ into the basket y_i (and leaving members of $\mathcal{L}_{K+1}$ at depth one) we arrive to a tree t whose basket nodes satisfy the Union condition (their light children do not violate the condition due to the assumption on $\mathcal{L}_i$, and their only heavy child having size $H+1$ does not violate the condition either since $\textsc{sumlight}(t, x) = \sum_{x \in \mathcal{L}_i} \textsc{size}(t_0, x) = H$ due also to the assumption on $\mathcal{L}_i$. Finally, the root node has K children (the initially empty baskets) of size $2H+2$ but since it has light children of total size H and a heavy child of size $H+1$, their sizes summing up to $2H+1$, the depth-one baskets also do not violate the Union condition. Hence $t_0 \vdash^* t$ for some Union tree t, thus is a Union-Find tree.
For the other direction, assume t_0 is a Union-Find tree. Then since t_0 is also a flat tree, it can be transformed into a Union tree t by repeatedly pushing nodes into baskets only. Since the pushed nodes' original parents are baskets as well (since they are parents of the basket into which we push), the set of nodes which are children of some basket do not change during the process. We also know that the charge of each node in the Union tree we gain at the end has to be zero, in particular, each basket x still has to have a heavy child of size exactly $H+1$ and $\textsc{sumlight}(t, x) = H$ has to hold. Let $y_1, \ldots, y_{K+1}$ stand for the basket nodes of t_0 (and of t as well): then, defining $\mathcal{L}_i$ as the set of light children of y_i in t suffices.

At this point we recall that the following problem 3 – Partition is **NP**-complete in the strong sense [5]: given a list $a_1, \ldots, a_{3m}$ of positive integers with the value $B = \frac{\sum_{i=1}^{3m} a_i}{m}$ being an integer, such that for each $1 \le i \le 3m$ we have $\frac{B}{4} < a_i < \frac{B}{2}$, does there exist a partition $\mathcal{B} = \{B_1, \ldots, B_k\}$ of the set $\{1, \ldots, 3m\}$ satisfying $\sum_{i \in B_j} a_i = B$ for each $1 \le j \le k$?
(Here "in the strong sense" means that the problem remains **NP**-complete even if the numbers are encoded in unary.)
Observe that by the condition $\frac{B}{4} < a_i < \frac{B}{2}$ each set B_j has to contain exactly three elements (the sum of any two of the input numbers is less than B and the sum of any four of them is larger than B), thus in particular in any solution above $k = m$ holds.
Note also that for any given instance $I = a_1, \ldots, a_{3m}$ of the above problem and offset $c \ge 0$, the instance $I' = a'_1, \ldots, a'_{3m}$ with $a'_i = a_i + c$ for each i and with $B' = B + 3c$, the set of solutions of I and I' coincide. Indeed, the instance I' still satisfies

$$\begin{aligned} \frac{B'}{4} &= \frac{B+3c}{4} = \frac{B}{4} + \frac{3}{4}c < \frac{B}{4} + c < a_i + c = a'_i \\ &= a_i + c < \frac{B}{2} + \frac{3}{2}c = \frac{B+3c}{2} = \frac{B'}{2}, \end{aligned}$$

hence each solution $\mathcal{B}' = \{B'_1, \ldots, B'_k\}$ of I' still contains triplets and $\sum_{i \in B'_j} a'_i = 3c + \sum_{i \in B'_j} a_i$ which is B' iff $\sum_{i \in B'_j} a_i = B$, thus any solution of I' is also a solution of I, the other direction being also straightforward to check.
Thus, by setting the above offset to $c = \lceil (1 + \max\{2^{\lceil \log B \rceil}, 2^4\} - B)/3 \rceil$ (in which case $B' = 2^D + d$ for some suitable integers $D > 3$ and $d \in \{1, 2, 3\}$) we get that the following problem is also strongly **NP**-complete:

Definition 1 (3 – Partition′). Input*: A sequence $a_1, \ldots, a_{3m}$ of positive integers such that $B = \frac{\sum_{i=1}^{3m} a_i}{m}$ is a positive integer of the form $2^D + d$ for some integer $D > 3$ and $d \in \{1, 2, 3\}$ and $\frac{B}{4} < a_i < \frac{B}{2}$ for each $i = 1, \ldots, 3m$.*
Output*: Does there exist a partition $\mathcal{B} = \{B_1, \ldots, B_k\}$ of the set $\{1, \ldots, 3m\}$ satisfying $\sum_{i \in B_j} a_i = B$ for each $j \in \{1, \ldots, k\}$?*
Observe that in any solution, $k = m$ and each B_j consists of exactly three elements.

We are now ready to show **NP**-completeness of recognition of Union-Find trees via giving a (logspace, many-to-one) reduction from the 3 – Partition′ problem to it. To each instance $I = a_1, \ldots, a_{3m}$ of the 3 – Partition′ problem we define the following flat tree $t(I)$:

- The number K of empty baskets in $t(I)$ is $m - 1$.
- The heaviness threshold parameter H of the tree is $B + 2^{D-1} - 1$ where $B = 2^D + d$ is the target sum $\frac{\sum_{i=1}^{3m} a_i}{m}$ with $D > 3$ being an integer and $d \in \{1, 2, 3\}$.
- There are $3m + (D-1)m$ light nodes of depth one in $t(I)$: first, to each member a_i of the input a light node x_i of size a_i is associated, and to each index $1 \le i \le m$ and $0 \le j < D - 1$, a light node y_i^j of size $w_i^j = 2^j$ is associated.

Note that since the 3 – Partition′ problem is strongly **NP**-complete we can assume that the input numbers a_i are encoded in unary, thus the tree $t(I)$ can be indeed built in logspace (and in polytime).
(See Figure 3 for an example.)
In order to ease notation we say that a member a_i of some sequence (or multiset) $a_1, \ldots, a_n$ satisfies the Union condition if the sum of the members of the sequence that are less than a_i is at least $a_i - 1$, i.e. $\sum_{a_j < a_i} a_j \ge a_i - 1$. In addition we say that the sequence itself satisfies the Union condition if each of its elements does so. It is clear that a node x of a tree t satisfies the Union condition if and only if the multiset $\{\textsc{size}(t, y) : y \in \textsc{children}(t, x)\}$ does so.
The following lemma states that the above construction is indeed a reduction:

LEMMA. For any instance I of the $3-$PARTITION$'$ problem, I has a solution iff $t(I)$ is a Union-Find tree.

PROOF. For one direction, assume $\{B_1, \ldots, B_m\}$ is a solution of the instance $I = a_1, \ldots, a_{3m}$. Then the multiset of the sizes of the light nodes can be partitioned into sets $\mathcal{L} = (L_1, \ldots, L_m)$ as $L_i = \{a_j : j \in B_i\} \cup \{w_i^j : 0 \leq j < D-1\}$. It is clear that $\sum_{\ell \in L_i} \ell = B + 2^{D-1} - 1 = H$. Thus, by Proposition 5 we only have to show that each L_i satisfies the Union condition. For the elements w_i^j of size 2^j it is clear since $\sum_{j'=0}^{j-1} w_i^{j'} = 2^j - 1$. Now let a be the smallest element of B_i. Then, since B_i consists of three integers summing up to B, we get that $a \leq \frac{B}{3}$. On the other hand, $\frac{B}{4} < a$ by the definition of the $3-$PARTITION$'$ problem. Recall that $B = 2^D + d$ for some $d \in \{1, 2, 3\}$: we get that $2^{D-2} < a$, thus in particular each weight w_i^j is smaller than a. Summing up all these weights we get $\sum_{j=0}^{D-2} w_i^j = 2^{D-1} - 1$. Applying $a \leq \frac{B}{3} = \frac{2^D+d}{3} < 2^{D-1} - 1$ we have that a satisfies the Union condition as well. Then, if $b > a$ is also a member of B_i, then it suffices to show $a + \sum_{j=0}^{D-2} w_i^j \geq b - 1$ but since $a \geq \frac{B}{4} > 2^{D-2}$ and $\sum_{j=0}^{D-2} w_i^j = 2^{D-1} - 1$, we get the left sum exceeds $2^{D-1} + 2^{D-2} - 1$ which is at least $b - 1$ since $b < \frac{B}{2} = \frac{2^D+d}{2} \leq 2^{D-1} + 2^{D-2}$ (since $d \leq 3$ and $2^{D-2} \geq 2$ if $D \geq 3$). Thus by Proposition 5, $t(I)$ is indeed a Union-Find tree. (See Figure 4.)

For the other direction, suppose $t(I)$ is a Union-Find tree. Then by Proposition 5, the multiset of the sizes of the light nodes can be partitioned into sets $\mathcal{L} = \{L_1, \ldots, L_m\}$ such that each L_i sums up exactly to $H = B + 2^{D-1} - 1$ and each L_i satisfies the Union condition.

First we show that each L_i contains exactly one "small" weight w_k^j for each $j = 0, \ldots, D-2$. (Note that each a_j exceeds 2^{D-2} hence the name of these weights.) We prove this by induction on j. Let a be the $j+1$th smallest integer in L_i. By induction, the lightest j members of L_i have sizes $2^0, 2^1, \ldots, 2^{j-1}$, summing up to $2^j - 1$. Since all the weights a_k have sizes more than 2^{D-2} as well as the weights $w_k^{j'}$ for $j' > j$, none of these can be the $j+1$th smallest integer in L_i without violating the Union condition. Thus, each set L_i has to have some $w_{k_i}^j$ as its $j+1$th smallest element, and since both the number of these sets and the number of the small weights of size 2^j is m we get that each set L_i contains exactly one small weight of size 2^j.

Thus, the small weights sum up exactly to $2^{D-1} - 1$ in each of the sets L_i, hence the weights a_i has to sum up exactly to $H - (2^{D-1} - 1) = 2^D + d = B$ in each of these sets, yielding a solution to the instance $(a_1, \ldots, a_{3m})$ of the $3-$PARTITION$'$ problem and concluding the proof.

Thus, since the strongly **NP**-complete $3-$PARTITION$'$ problem reduces to the problem of *flat* trees being Union-Find trees, and moreover, if a flat tree is a Union-Find tree, then it can be constructed from a Union tree by applying finitely many path compressions (all one has to do is to "move out" the light nodes from the baskets by calling a find operation on each of them successively), we have proven the following:

Theorem 3. *It is* **NP***-complete already for flat trees to decide whether*

i) a given (flat) tree is a Union-Find tree and

ii) whether a given (flat) tree can be constructed from a Union tree by applying a number of path compression operations.

5. Conclusion, future directions

We have shown that unless **P** = **NP**, there is no efficient algorithm to check whether a given tree is a valid Union-Find tree, assuming union-by-size strategy and usage of path compression since the problem is **NP**-complete. A natural question is whether the recognition problem remains so hard assuming union-by-rank strategy (and of course path compression)? Since the heights of merged trees does not add up, only increases by at most one, there is no "obvious" way to encode arithmetic into the construction of these trees, and even the characterization by the push operation is unclear to hold in that case (since in that setting, path compression can alter the order of the subsequent merging).

References

References

[1] Leizhen Cai. The recognition of union trees. *Inf. Process. Lett.*, 45(6):279–283, 1993.

[2] Thomas H. Cormen, Clifford Stein, Ronald L. Rivest, and Charles E. Leiserson. *Introduction to Algorithms*. McGraw-Hill Higher Education, 2nd edition, 2001.

[3] M. Fredman and M. Saks. The cell probe complexity of dynamic data structures. In *Proceedings of the Twenty-first Annual ACM Symposium on Theory of Computing*, STOC '89, pages 345–354, New York, NY, USA, 1989. ACM.

[4] Bernard A. Galler and Michael J. Fisher. An improved equivalence algorithm. *Commun. ACM*, 7(5):301–303, May 1964.

[5] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.

[6] Kevin Knight. Unification: A multidisciplinary survey. *ACM Comput. Surv.*, 21(1):93–124, March 1989.

[7] Robert Endre Tarjan. Efficiency of a good but not linear set union algorithm. *J. ACM*, 22(2):215–225, April 1975.

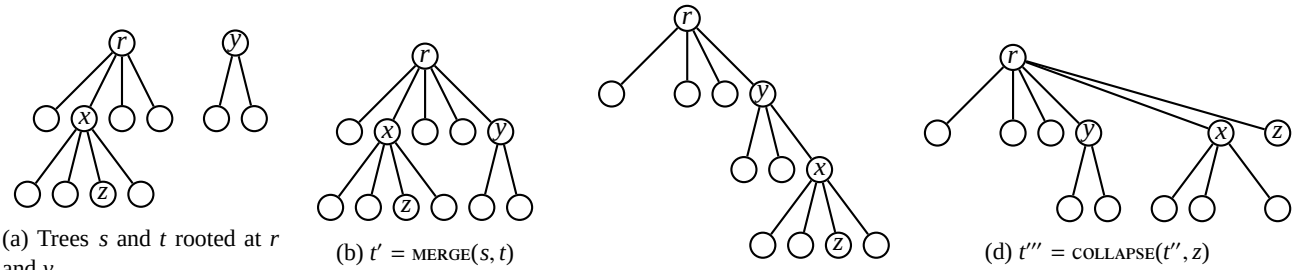

(a) Trees s and t rooted at r and y

(b) $t' = \text{MERGE}(s, t)$

(c) $t'' = \text{PUSH}(t', x, y)$

(d) $t''' = \text{COLLAPSE}(t'', z)$

Figure 1: Merge, collapse and push.

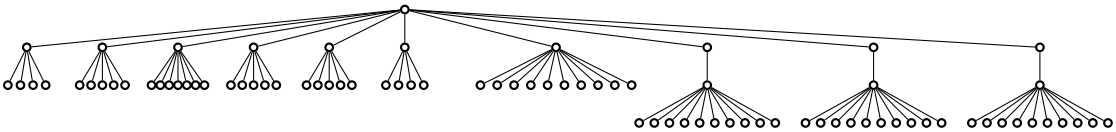

Figure 2: A flat tree with $H = 9$, $K = 3$. The six depth-one light nodes' sizes sum up to $4 \times 9 = 36$.

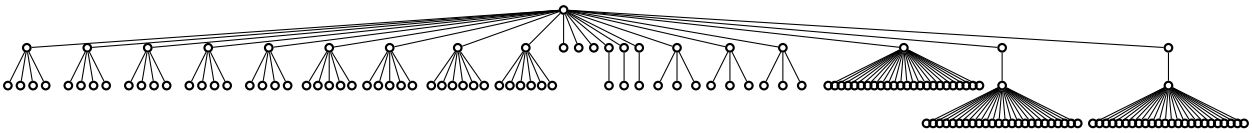

Figure 3: Illustrating the reduction: $t(I)$ for $I = (5, 5, 5, 5, 5, 6, 6, 7, 7)$, $m = 3$, $B = 17$. Then $D = 4$, $d = 1$ and $H = 24$. Small weights w_i^j are of size $1, 2, 4$.

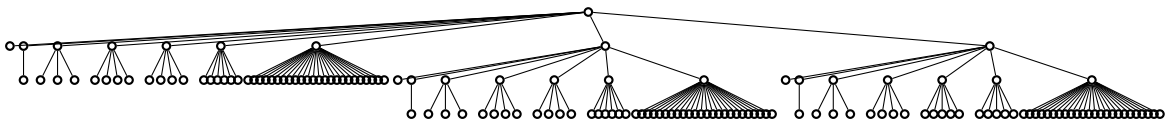

Figure 4: The Union tree corresponding to a solution of the instance on Figure 3.